

Perfectly Secure Message Transmission in Two Rounds

Gabriele Spini ^{*1,2,3} and Gilles Zémor ^{†1}

¹Institut de Mathématiques de Bordeaux, UMR 5251, Université de Bordeaux

²Mathematical Institute, Leiden University

³CWI Amsterdam

Abstract

In the model that has become known as “Perfectly Secure Message Transmission” (PSMT), a sender Alice is connected to a receiver Bob through n parallel two-way channels. A computationally unbounded adversary Eve controls t of these channels, meaning she can acquire and alter any data that is transmitted over these channels. The sender Alice wishes to communicate a secret message to Bob *privately* and *reliably*, i.e. in such a way that Eve will not get any information about the message while Bob will be able to recover it completely.

In this paper, we focus on protocols that work in two transmission rounds for $n = 2t + 1$. We break from previous work by following a conceptually simpler blueprint for achieving a PSMT protocol. We reduce the previously best-known communication complexity, i.e. the number of transmitted bits necessary to communicate a 1-bit secret, from $O(n^3 \log n)$ to $O(n^2 \log n)$. Our protocol also answers a question raised by Kurosawa and Suzuki and hitherto left open: their protocol reaches optimal transmission rate for a secret of size $O(n^2 \log n)$ bits, and the authors raised the problem of lowering this threshold. The present solution does this for a secret of $O(n \log n)$ bits. Additionally, we show how our protocol can be adapted to a Network Coding context.

Keywords: Perfectly Secure Message Transmission

1 Introduction

The problem of Perfectly Secure Message Transmission (PSMT for short) was introduced by Dolev et al. in [4] and involves two parties, a sender Alice and a receiver Bob, who communicate over n parallel channels in the presence of an adversary Eve. Eve is computationally unbounded and controls $t \leq n$ of the channels, meaning that she can read and overwrite any data sent over the channels under her control. The goal of PSMT is to design a protocol that allows Alice to communicate a secret message to Bob *privately* and *reliably*, i.e. in such a way that Eve will not be able to acquire any information on the message, while Bob will always be able to completely recover it.

^{*}Supported by the Algant-Doc doctoral program, www.algant.eu

[†]Supported by the “Investments for the Future” Programme IdEx Bordeaux – CPU (ANR-10-IDEX- 03-02).

A shorter version of this article, under the same title, © IACR, is scheduled to appear in the Proceedings of the Theory of Cryptography Conference (TCC) 2016-B.

Two factors influence whether PSMT is possible and how difficult it is to achieve, namely the number t of channels corrupted and controlled by Eve, and the number r of transmission rounds, where a transmission round is a phase involving only one-way communication (either from Alice to Bob, or from Bob to Alice).

It was shown in Dolev et al.'s original paper [4] that for $r = 1$, i.e. when communication is only allowed from Alice to Bob, PSMT is possible if and only if $n \geq 3t + 1$. It was also shown in [4] that for $r \geq 2$, i.e. when communication can be performed in two or more rounds, PSMT is possible if and only if $n \geq 2t + 1$, although only a very inefficient way to do this was proposed. A number of subsequent efforts were made to improve PSMT protocols, notably in the most difficult case, namely for two rounds and when $n = 2t + 1$. The following two quantities, called *communication complexity* and *transmission rate*, were introduced and give a good measure of the efficiency of a PSMT protocol. They are defined as follows:

Communication complexity := total number of bits transmitted to
communicate a single-bit secret,

Transmission rate := $\frac{\text{total number of bits transmitted}}{\text{bit-size of the secret}}.$

Focusing exclusively on the case $n = 2t + 1$, Dolev et al. [4] presented a PSMT protocol for $r = 3$ with transmission rate $O(n^5)$: for $r = 2$ a protocol was presented with non-polynomial rate.

Sayeed and Abu-Amara [15] were the first to propose a two-round protocol with a polynomial transmission rate of $O(n^3)$. They also achieved communication complexity of $O(n^3 \log n)$. Further work by Agarwal et al. [1] improved the transmission rate to $O(n)$ meeting, up to a multiplicative constant, the lower bound of [20]. However, this involved exponential-time algorithms for the participants in the protocol. The current state-of-the art protocol is due to Kurosawa and Suzuki [9, 10]; it achieves $O(n)$ transmission rate with a polynomial-time effort from the participants. All these protocols do not do better than $O(n^3 \log n)$ for the communication complexity.

We contribute to this topic in the following ways. We present a constructive protocol for which only polynomial-time, straightforward computations are required of the participants, that achieves the improved communication complexity of $O(n^2 \log n)$. In passing, we give an affirmative answer to an open problem of Kurosawa and Suzuki (at the end of their paper [10]) that asks whether it is possible to achieve the optimal transmission rate $O(n)$ for a secret of size less than $O(n^2 \log n)$ bits. We do this for a secret of $O(n \log n)$ bits.

Just as importantly, our solution is conceptually significantly simpler than previous protocols. Two-round PSMT involves Bob initiating the protocol by first sending an array of symbols (x_{ij}) over the n parallel channels, where the first index i means that symbol x_{ij} is sent over the i -th channel. All previous proposals relied on arrays (x_{ij}) with a lot of structure, with linear relations between symbols that run both along horizontal (constant j) and vertical (constant i) lines. In contrast, we work with an array (x_{ij}) consisting of completely independent rows $\mathbf{x}^{(j)} = (x_{1j}, x_{2j}, \dots, x_{nj})$ that are simply randomly chosen words of a given Reed-Solomon code. In its simplest, non-optimized, form, the PSMT protocol we present only involves simple syndrome computations from Alice, and one-time padding the secrets it wishes to transfer with the image of linear forms applied to corrupted versions of the codewords $\mathbf{x}^{(j)}$ it has received from Bob.

Arguably, the method could find its way into textbooks as a relatively straightforward application of either secret-sharing or wiretap coset-coding techniques. In its optimized form, the protocol retains sufficient simplicity to achieve a transmission rate $5n + o(n)$, compared to the previous

record of $6n + o(n)$ of [6] obtained by painstakingly optimizing the $25n + o(n)$ transmission rate of [10].

Finally, we show how our protocol can be adapted to provide security in a more general setting, where Eve can eavesdrop linear combination of the transmitted symbols and inject linear combinations of errors from a set of her choice. This means that thanks to its simpler core structure, our protocol has greater potential of being adapted to more complex scenarios.

In the next Section we give an overview of our method and techniques.

2 Protocol Overview

The procedure takes as input the number $n = 2t + 1$ of channels between Alice and Bob and the number ℓ of secret messages to be communicated; we assume that the messages lie in a finite field \mathbb{F}_q . First, a code \mathcal{C} that will be the basic communication tool is selected; \mathcal{C} is a linear block code of length n over \mathbb{F}_q , dimension $t + 1$ and minimum distance $t + 1$. It furthermore has the property that the knowledge of t symbols of any of its codewords \mathbf{x} leaves $\mathbf{h}\mathbf{x}^T$ completely undetermined, where \mathbf{h} is a vector produced together with \mathcal{C} at the beginning of the protocol. The code \mathcal{C} can be a Reed-Solomon code.

Since we require at most two rounds of communication, Bob starts the procedure; he chooses a certain number of random and independent codewords \mathbf{x} , and communicates them by sending the i -th symbol of each codeword over the i -th channel. This is a first major difference from previous papers, notably [10], where codewords are communicated in a more complicated “horizontal-and-vertical” fashion; our construction is thus conceptually simpler and eliminates techniques introduced by early papers [15] which marked substantial progress at the time but also hindered the development of more efficient protocols when they survived in subsequent work.

As a result of this first round of communication, Alice receives a corrupted version $\mathbf{y} = \mathbf{x} + \mathbf{e}$ for each codeword \mathbf{x} sent by Bob. As in previous PSMT protocols, Alice then proceeds by broadcast, meaning every symbol she physically sends to Bob, she sends n times, once over every channel i . In this way privacy is sacrificed, since Eve can read everything Alice sends, but reliability is ensured, since Bob recovers every transmitted symbol by majority decoding.

A secret message consisting of a single symbol $s \in \mathbb{F}_q$ is encoded by Alice as $s + \mathbf{h}\mathbf{y}^T$ for some received vector \mathbf{y} . In other words, s is one-time padded with the quantity $\mathbf{h}\mathbf{y}^T$ and this is broadcast to Bob. Notice that at this point, revealing $s + \mathbf{h}\mathbf{y}^T$ to Eve gives her zero information on s . This is because she can have intercepted at most t symbols of the codeword \mathbf{x} : therefore the element $\mathbf{h}\mathbf{x}^T$ is completely unknown to her by the above property of \mathcal{C} and \mathbf{h} , and the mask $\mathbf{h}\mathbf{y}^T = \mathbf{h}\mathbf{x}^T + \mathbf{h}\mathbf{e}^T$ is unknown to her as well.

Now broadcasting the quantity $s + \mathbf{h}\mathbf{y}^T$ is not enough by itself to convey the secret s to Bob, because Bob also does not have enough information to recover the mask $\mathbf{h}\mathbf{y}^T$. To make the protocol work, Alice needs to give Bob extra information that tells Eve nothing she doesn’t already know.

This extra information comes in two parts. The first part is simply the syndrome $\sigma(\mathbf{y}) = \mathbf{H}\mathbf{y}^T$ of \mathbf{y} , where \mathbf{H} is a parity-check matrix of \mathcal{C} ; notice that this data is indeed useless to Eve, who already knows it given that $\mathbf{H}\mathbf{y}^T = \mathbf{H}\mathbf{x}^T + \mathbf{H}\mathbf{e}^T = \mathbf{H}\mathbf{e}^T$ where \mathbf{e} is chosen by herself.

The second part makes use of the fact that during the first phase, Bob has not sent a single codeword \mathbf{x} to Alice, but a batch of codewords \mathcal{X} and Alice has received a set \mathcal{Y} of vectors made up of the corrupted versions $\mathbf{y} = \mathbf{x} + \mathbf{e}$ of the codewords \mathbf{x} . Alice will sacrifice a chosen subset

of these vectors \mathbf{y} and reveal them completely to Bob and Eve by broadcast. Note that this does not yield any information on the unrevealed vectors \mathbf{y} since Bob has chosen the codewords \mathbf{x} of \mathcal{X} randomly and independently. At this point we apply an idea that originates in [10]: the chosen revealed subset of \mathcal{Y} is called in [10] a *pseudo-basis* of \mathcal{Y} . To compute a pseudo-basis of \mathcal{Y} , Alice simply computes all syndromes $\sigma(\mathbf{y})$ for $\mathbf{y} \in \mathcal{Y}$, and chooses a minimal subset of \mathcal{Y} whose syndromes generate linearly all syndromes $\sigma(\mathbf{y})$ for $\mathbf{y} \in \mathcal{Y}$. A pseudo-basis of \mathcal{Y} could alternatively be called a *syndrome-spanning* subset of \mathcal{Y} . Now elementary coding-theory arguments imply that the syndrome function σ is injective on the subspace generated by the set of *all* errors \mathbf{e} that Eve applies to all Bob's codewords \mathbf{x} (Lemma 2 and Proposition 1). Therefore a pseudo-basis of \mathcal{Y} gives Bob access to the whole space spanned by Eve's errors and allows him, *for any non-revealed* $\mathbf{y} = \mathbf{x} + \mathbf{e}$, to recover the error \mathbf{e} from the syndrome $\sigma(\mathbf{y}) = \sigma(\mathbf{e})$.

The above protocol is arguably “the right way” of exploiting the pseudo-basis idea of Kurosawa and Suzuki, by which we mean it is the simplest way of turning it into a two-round PSMT protocol. We shall present optimized variants that achieve the communication complexity and transmission rate claimed in the Introduction. Our final protocol involves two additional ideas; the first involves a more efficient broadcasting scheme than pure repetition: this idea was also used by Kurosawa and Suzuki. The second idea is new and involves using a decoding algorithm for the code \mathcal{C} .

The rest of the paper is organized as follows. In Section 3 we recall the coding theory that we need to set up the protocol. In particular, Section 3.1 introduces the code \mathcal{C} and the vector \mathbf{h} with the desired properties. Section 3.2 introduces Kurosawa and Suzuki's pseudo-basis idea, though we depart somewhat from their original description to fit our syndrome-coding approach to PSMT.

In Section 4 we describe in a formal way the protocol sketched above, and we compute its communication cost; it will turn out that this construction has a communication complexity of $O(n^3 \log n)$ and a transmission rate of $O(n^2)$.

Section 5 is devoted to improving the efficiency of the protocol; specifically, Section 5.1 introduces generalized broadcast, Sections 5.2 and 5.3 show how to lower the cost of transmitting the pseudo-basis, while Section 5.4 presents a way to improve the efficiency of the last part of the protocol. A key aspect of this section is that Alice must make extensive use of a *decoding* algorithm for linear codes, a new feature compared to previous work on the topic.

In Section 6 we implement these improvements and compute the cost of the resulting protocol, reaching a communication complexity of $O(n^2 \log n)$ and a transfer rate of $5n + o(n)$. We also show in this section that optimal transfer rate is achieved for a secret of $O(n \log n)$ bits.

Finally, in Section 7 we adapt the ideas developed for classical PSMT to a generalized setting with a more powerful eavesdropper. This has applications to achieving security in a Network Coding context. Section 8 gives concluding remarks.

3 Setting and Techniques

3.1 Error-Correcting Codes for Communication

We will use the language of Coding Theory, for background, see e.g. [13]. Let us briefly recall that when a linear code over the finite field \mathbb{F}_q is defined as $\mathcal{C} = \{\mathbf{x} \in \mathbb{F}_q^n, \mathbf{H}\mathbf{x}^T = 0\}$, the $r \times n$ matrix

\mathbf{H} is called a *parity-check matrix* for \mathcal{C} and the mapping

$$\begin{aligned}\sigma : \mathbb{F}_q^n &\rightarrow \mathbb{F}_q^r \\ \mathbf{x} &\mapsto \mathbf{H}\mathbf{x}^T\end{aligned}$$

is referred to as the *syndrome map*. Recall also that a code of parameters (length, dimension, minimum Hamming distance) $[n, k, d]$ is said to be Maximum Distance Separable or *MDS*, if $d + k = n + 1$. Particular instances of MDS codes are Reed-Solomon codes, which exist whenever the field size q is equal to or larger than the length n . In a secret-sharing context, Reed-Solomon codes are equivalent to Shamir's secret-sharing scheme [16], and they have been used extensively to construct PSMT protocols. We could work from the start with Reed-Solomon codes, equivalently Shamir's scheme, but prefer to use more general MDS codes, not purely for generality's sake, but to stay unencumbered by polynomial evaluations and to highlight that we have no need for anything other than Hamming distance properties. In section 6, we will need our MDS codes to come with a decoding algorithm and will have to invoke Reed-Solomon codes specifically: we will only need to know of the existence of a polynomial-time algorithm though, and will not require knowledge of any specifics.

We will need an MDS code \mathcal{C} that will be used to share randomness, together with a vector \mathbf{h} such that the value of $\mathbf{h}\mathbf{x}^T$ is completely undetermined for a codeword $\mathbf{x} \in \mathcal{C}$ even when t symbols of \mathbf{x} are known. The linear combination given by \mathbf{h} will then be used to create the masks that hide the secrets.

The following Lemma states the existence of such a pair $(\mathcal{C}, \mathbf{h})$: it is a slightly non-standard use of Massey's secret sharing scheme [14]. It is implicit that we suppose $q > n$, so that MDS codes exist for all dimensions and length up to $n + 1$.

Lemma 1. *For any n and any $t < n$ there exists an MDS code \mathcal{C} of parameters $[n, t + 1, n - t]$ and a vector $\mathbf{h} \in \mathbb{F}_q^n$ such that given a random codeword $\mathbf{x} \in \mathcal{C}$, the scalar product $\mathbf{h}\mathbf{x}^T$ is completely undetermined even when t symbols of \mathbf{x} are known.*

Proof. Let \mathcal{C}' be an MDS code of parameters $[n + 1, t + 1, n - t + 1]$; notice that such a code exists for any n and $t \leq n$ [13]. Let \mathcal{C} be the code obtained from \mathcal{C}' by puncturing at its last coordinate, i.e.

$$\mathcal{C} := \{\mathbf{x} \in \mathbb{F}_q^n : \exists x \in \mathbb{F}_q \text{ with } (\mathbf{x}, x) \in \mathcal{C}'\}$$

The minimum distance of \mathcal{C} is at most one less than that of \mathcal{C}' , and \mathcal{C} is MDS of parameters $[n, t + 1, n - t]$ as requested. Now let \mathbf{H}' be a parity-check matrix of \mathcal{C}' ; since \mathcal{C}' has minimum distance $n - t + 1 > 1$, there is at least one row of \mathbf{H}' whose last symbol is non-zero, i.e. such a row is of the form

$$(\mathbf{h}, \alpha) \in \mathbb{F}_q^{n+1} \text{ with } \mathbf{h} \in \mathbb{F}_q^n, 0 \neq \alpha \in \mathbb{F}_q.$$

We claim that the pair $(\mathcal{C}, \mathbf{h})$ is of the desired type: indeed, let \mathbf{x} be a random codeword of \mathcal{C} . Then there exists a (unique) codeword \mathbf{x}' of \mathcal{C}' such that $\mathbf{x}' = (\mathbf{x}, x)$; now $\mathbf{h}\mathbf{x}^T = -\alpha x$, i.e. the knowledge of $\mathbf{h}\mathbf{x}^T$ is equivalent to the knowledge of x , given that α is non-zero.

Now since \mathcal{C}' has dimension $t + 1$, for any t known symbols x_{i_1}, \dots, x_{i_t} of \mathbf{x} and any $\tilde{x} \in \mathbb{F}_q$, there exists exactly one $\tilde{\mathbf{x}}' \in \mathcal{C}'$ such that $\tilde{\mathbf{x}}'_{i_j} = x_{i_j}$ for any j and such that $\tilde{\mathbf{x}}'_{n+1} = \tilde{x}$. Hence the claim holds. □

As stated above, this lemma will guarantee the privacy of our protocols; conversely, we can achieve reliable (although not private) communication via the following remark: Alice and Bob can *broadcast* a symbol by sending it over all the channels; since Eve only controls $t < n/2$ of them, the receiver will be able to correct any error introduced by Eve with a simple majority choice. Broadcast thus guarantees reliability by sacrificing privacy.

3.2 Pseudo-Bases or Syndrome-Spanning Subsets

The second fundamental building block of our paper is the notion of *pseudo-basis*, introduced by Kurosawa and Suzuki [10]. The concept stems from the following intuition: assume that Bob communicates a single codeword \mathbf{x} of an MDS code \mathcal{C} to Alice by sending each of its n symbols over the corresponding channel. Eve intercepts t of these symbols, thus \mathcal{C} must have dimension at least $t + 1$ if we want to prevent her from learning \mathbf{x} ; but this means that the minimum distance of \mathcal{C} cannot exceed $n + 1 - (t + 1) = t + 1$, which is not enough for Alice to correct an arbitrary pattern of up to t errors that Eve can introduce.

If, however, we repeat the process for several different $\mathbf{x}^{(i)}$, then Alice and Bob have an important advantage: they know that all the errors introduced by Eve *always lie in the same subset of t coordinates*. Kurosawa and Suzuki propose the following strategy to exploit this knowledge: Alice can compute a pseudo-basis (a subset with special properties) of the received vectors; she can then transmit it to Bob, who will use this special structure of the errors to determine their support.

The key is the following simple lemma:

Lemma 2. *Let \mathcal{C} be a linear code of parameters $[n, k, d]_q$, and let \mathbf{H} be a parity-check matrix of \mathcal{C} ; let E be a linear subspace of vectors of \mathbb{F}_q^n such that the Hamming weight $w_H(\mathbf{e})$ of \mathbf{e} satisfies $w_H(\mathbf{e}) < d$ for any $\mathbf{e} \in E$.*

We then have that the following map is injective:

$$\begin{aligned} \sigma|_E : E &\rightarrow \mathbb{F}_q^{n-k} \\ \mathbf{e} &\mapsto \mathbf{H}\mathbf{e}^T \end{aligned}$$

Proof. Simply notice that $\ker(\sigma|_E) = \{\mathbf{0}\}$: indeed, $\ker(\sigma|_E) \subseteq \mathcal{C}$; but by assumption all elements of E have weight smaller than d , so that $\ker(\sigma|_E) = \{\mathbf{0}\}$. □

We can now introduce the concept of pseudo-basis; for the rest of this section, we assume that a linear code \mathcal{C} of parameters $[n, k, d]_q$ has been chosen, together with a parity-check matrix \mathbf{H} and associated syndrome map σ .

Definition 1 (Pseudo-Basis [10]). Let \mathcal{Y} be a set of vectors of \mathbb{F}_q^n ; a *pseudo-basis* of \mathcal{Y} is a subset $\mathcal{W} \subseteq \mathcal{Y}$ such that $\sigma(\mathcal{W})$ is a basis of the syndrome subspace $\langle \sigma(\mathcal{Y}) \rangle$.

Notice that a pseudo-basis has thus cardinality at most $n - k$, and that it can be computed in time polynomial in n .

The following property formalizes the data that Bob can acquire after he obtains a pseudo-basis of the words received by Alice:

Proposition 1 ([10]). *Let $\mathcal{X}, \mathcal{E}, \mathcal{Y}$ be three subsets:*

$$\begin{aligned}\mathcal{X} &:= \{\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(r)}\} \subseteq \mathcal{C}, \\ \mathcal{E} &:= \{\mathbf{e}^{(1)}, \dots, \mathbf{e}^{(r)}\} \subseteq \mathbb{F}_q^n \quad \text{such that } \#\bigcup (\text{support}(\mathbf{e}^{(j)}) : j = 1, \dots, r) < d, \\ \mathcal{Y} &:= \{\mathbf{y}^{(1)}, \dots, \mathbf{y}^{(r)}\} \subseteq \mathbb{F}_q^n \quad \text{with } \mathbf{y}^{(j)} = \mathbf{x}^{(j)} + \mathbf{e}^{(j)} \text{ for every } j\end{aligned}$$

Then, given knowledge of \mathcal{X} and a pseudo-basis of \mathcal{Y} , we can compute $\mathbf{e}^{(j)}$ from its syndrome $\sigma(\mathbf{e}^{(j)})$, for any $1 \leq j \leq r$.

Proof. The hypothesis on the supports of the elements of \mathcal{E} implies that the subspace $E = \langle \mathcal{E} \rangle$ satisfies the hypothesis of Lemma 2 and the syndrome function is therefore injective on $\langle \mathcal{E} \rangle$. Given the pseudo-basis $\{\mathbf{y}^{(i)} : i \in I\}$, we can decompose any syndrome $\sigma(\mathbf{e}^{(j)})$ as

$$\begin{aligned}\sigma(\mathbf{e}^{(j)}) &= \sum_{i \in I} \lambda_i \sigma(\mathbf{y}^{(i)}) = \sum_{i \in I} \lambda_i \sigma(\mathbf{e}^{(i)}) \\ &= \sigma \left(\sum_{i \in I} \lambda_i \mathbf{e}^{(i)} \right)\end{aligned}$$

which yields

$$\mathbf{e}^{(j)} = \sum_{i \in I} \lambda_i \mathbf{e}^{(i)}$$

by injectivity of σ on E (cf. Lemma 2). □

Remark 1. Since the syndrome map induces a one-to-one mapping from E to $\sigma(E)$, we also have that $\{\mathbf{y}^{(i)} : i \in I\}$ is a pseudo-basis of \mathcal{Y} if and only if $\{\mathbf{e}^{(i)} : i \in I\}$ is a basis of $E = \langle \mathcal{E} \rangle$.

The reader should now have a clear picture of how the pseudo-basis will be used to obtain shared randomness: Bob will select a few codewords $\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(r)}$ in an MDS code of distance at least $t+1$, then communicate them to Alice by sending the i -th symbol of each codeword over channel i ; Alice will be able to compute a pseudo-basis of the received words, a clearly non-expensive computation, then communicate it to Bob. Bob will then be able to determine any error introduced by Eve just from its syndrome as just showed in Proposition 1.

The following section gives all the details.

4 A First Protocol

We now present the complete version of our first communication protocol, following the blueprint of Section 2.

Protocol 1. The protocol allows Alice to communicate ℓ secret elements $s^{(1)}, \dots, s^{(\ell)}$ of \mathbb{F}_q to Bob, where q is an arbitrary integer with $q > n$. The protocol takes as input an MDS code \mathcal{C} of parameters $[n, t+1, t+1]_q$ and a vector \mathbf{h} of length n as in Lemma 1.

- I. Bob chooses $t + \ell$ uniformly random and independent codewords $\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(t+\ell)}$ of \mathcal{C} and communicates them to Alice by sending the i -th symbol of each codeword over the i -th channel.
- II. Alice receives the corrupted versions $\mathbf{y}^{(1)} = \mathbf{x}^{(1)} + \mathbf{e}^{(1)}, \dots, \mathbf{y}^{(t+\ell)} = \mathbf{x}^{(t+\ell)} + \mathbf{e}^{(t+\ell)}$; she then proceeds with the following actions:
 - (i) She computes a pseudo-basis $\{\mathbf{y}^{(i)} : i \in I\}$ for $I \subset \{1, \dots, t + \ell\}$ of the received values and broadcasts to Bob $(i, \mathbf{y}^{(i)} : i \in I)$.
 - (ii) She then considers the first ℓ words that do *not* belong to the pseudo-basis; to ease the notation, we will re-name them $\mathbf{y}^{(1)}, \dots, \mathbf{y}^{(\ell)}$. For each secret $s^{(j)}$ to be communicated she broadcasts to Bob the following two elements:
 - $\mathbf{H}(\mathbf{y}^{(j)})^T$, the syndrome of $\mathbf{y}^{(j)}$;
 - $s^{(j)} + \mathbf{h}(\mathbf{y}^{(j)})^T$.
- III. Proposition 1 guarantees that for any j , $1 \leq j \leq \ell$, Bob can compute the error vector $\mathbf{e}^{(j)}$ and hence reconstruct $\mathbf{y}^{(j)} = \mathbf{x}^{(j)} + \mathbf{e}^{(j)}$ from his knowledge of $\mathbf{x}^{(j)}$. He can therefore open the mask $\mathbf{h}(\mathbf{y}^{(j)})^T$ and obtain the secret $s^{(j)}$.

Proposition 2. *The above protocol allows for private and reliable communication of ℓ elements of \mathbb{F}_q .*

Proof. As a first remark, notice that since the pseudo-basis has cardinality at most t as remarked in Definition 1, Alice has enough words to mask her ℓ secret messages, since the total number of words is equal to $t + \ell$. We can now prove that the protocol is private and reliable:

- *Privacy:* Eve can intercept at most t coordinates of each codeword sent over the channels in the first step; the codewords corresponding to the pseudo-basis are revealed in step II-(i), but this information is useless since the words are chosen independently and those belonging to the pseudo-basis are no longer used. For any $\mathbf{y}^{(j)}$ that does not belong to the pseudo-basis, the syndrome $\mathbf{H}(\mathbf{y}^{(j)})^T$ is also transmitted, but Eve already knows it since $\mathbf{H}(\mathbf{y}^{(j)})^T = \mathbf{H}(\mathbf{x}^{(j)} + \mathbf{e}^{(j)})^T = \mathbf{H}(\mathbf{e}^{(j)})^T$, where $\mathbf{e}^{(j)}$ denotes the error she introduced herself on $\mathbf{x}^{(j)}$.

Hence thanks to Lemma 1, Eve has no information on any $\mathbf{h}(\mathbf{y}^{(j)})^T$, so that privacy holds.

- *Reliability:* Eve can disrupt the communication only at step I, since all the following ones only use broadcasts. Proposition 1 then ensures that Bob can recover the vectors $\mathbf{y}^{(j)}$ from their syndromes and the corresponding codeword $\mathbf{x}^{(j)}$. From there he can compute and remove the mask $\mathbf{h}(\mathbf{y}^{(j)})^T$ without error.

□

We now compute the communication complexity and transmission rate of this first protocol, underlining the most expensive parts:

Communication complexity: we can set $\ell := 1$.

- Step I requires transmitting $t+1$ codewords over the channels, thus requiring a total of $O(n^2)$ symbols to be transmitted.
- Step II-(i) requires broadcasting up to t words of \mathbb{F}_q^n , thus giving a total of $O(n^3)$ symbols to be transmitted.
- Finally, step II-(ii) requires broadcasting a total of $t+1$ symbols (a size- t syndrome and the masked secret), thus giving a total of $O(n^2)$ elements to be transmitted.

Hence since we can assume that $q = O(n)$, we get a total communication complexity of

$$O(n^3 \log n)$$

bits to be transmitted to communicate a single-bit secret.

Transfer rate: optimal rate is achieved for $\ell = \Omega(n)$.

- Step I requires transmitting $t + \ell$ codewords, for a total of $O(n^2 + n\ell)$ symbols.
- Step II-(i) remains unchanged from the single-bit case, and thus requires transmitting $O(n^3)$ symbols.
- Finally, step II-(ii) requires broadcasting a total of $\ell(t+1)$ symbols (ℓ size- t syndromes and the masked secrets), thus giving a total of $O(n^2\ell)$ symbols;

To sum up, the overall transmission rate is equal to

$$\frac{O(n^2 + n\ell + n^3 + n^2\ell)}{\ell} = O(n^2).$$

It is immediately seen that the main bottleneck for communication complexity is step II-(i), i.e. the communication of the pseudo-basis, while for transmission rate it is step II-(ii), i.e. the communication of the masked secrets and of the syndromes. We address these issues in the following sections.

5 Improvements to the Protocol

We discuss in this section some key improvements to the protocol; Section 5.1 presents the key technique of generalized broadcast, Sections 5.2 and 5.3 show a new way to communicate the pseudo-basis (the main bottleneck for communication complexity) and Section 5.4 a new way to communicate the masked secret and the information to open the masks (bottleneck for transmission rate).

5.1 Generalized Broadcast

Our improvements on the two bottlenecks showed in Section 4 rely on the fundamental technique of generalized broadcast, which has been highlighted in the paper by Kurosawa and Suzuki [10].

The intuition is the following: we want to choose a suitable code $\mathcal{C}_{\text{BCAST}}$ for perfectly reliable transmission, i.e. we require that if any word $\mathbf{x} \in \mathcal{C}_{\text{BCAST}}$ is communicated by sending each symbol \mathbf{x}_i over the i -th channel, then \mathbf{x} can always be recovered in spite of the errors introduced by the adversary. In the general situation, since Eve can introduce up to t errors, $\mathcal{C}_{\text{BCAST}}$ must have minimum distance $2t + 1 = n$, and hence dimension 1; for instance, $\mathcal{C}_{\text{BCAST}}$ can be a repetition code, yielding the broadcast protocol of Section 3.1.

Now assume that at a certain point of the protocol, Bob gets to know the position of m channels under Eve's control; then the communication system between the two has been improved: instead of n channels with t errors, we have n channels with m erasures and $t - m$ errors (since Bob can ignore the symbols received on the m channels under Eve's control that he has identified). We can thus expect that reliable communication between Alice and Bob (i.e., broadcast) can be performed at a lower cost by using a code with smaller distance and greater dimension; the following lemma formalizes this intuition.

Lemma 3 (Generalized Broadcast). *Let $m \leq t$ and let \mathcal{C}_m be an MDS code of parameters $[n, m + 1, n - m]_q$; assume that Bob knows the location of m channels controlled by Eve. Then Alice can communicate with perfect reliability $m + 1$ symbols x_1, \dots, x_{m+1} of \mathbb{F}_q to Bob in the following way: she first takes the codeword $\mathbf{c} \in \mathcal{C}_m$ which encodes (x_1, \dots, x_{m+1}) , then sends each symbol of \mathbf{c} through the corresponding channel; Eve cannot prevent Bob from completely recovering the message.*

We refer to this procedure as m -generalized broadcast.

Proof. Notice that \mathbf{c} is well-defined since \mathcal{C}_m has dimension $m + 1$. Now since Bob knows the location of m channels that are under Eve's control, he can replace the symbols of \mathbf{c} received via these channels with erasure marks \perp , and consider the truncated codeword $\tilde{\mathbf{c}}$ lacking these symbols. Now $\tilde{\mathbf{c}}$ belongs to the punctured code obtained from \mathcal{C}_m by removing m coordinates, which has minimum distance $(n - m) - m \geq 2(t - m) + 1$; it can thus correct up to $t - m$ errors, which is exactly the maximum number of errors that Eve can introduce (since she controls at most $t - m$ of the remaining channels). Once he has obtained the shortened codeword $\tilde{\mathbf{c}}$, he can then recover the complete one since \mathcal{C}_m can correct from m erasures, given that it has minimum distance $n - m \geq m$. \square

Hence if Alice knows that Bob has identified at least m channels under Eve's control, she can divide the cost of a broadcast by a factor m (since the above method requires to transmit n symbols of \mathbb{F}_q to communicate $m + 1$ symbols of \mathbb{F}_q).

In the following sections we will make use of Lemma 3 to improve the efficiency of the protocol.

5.2 Improved Transmission of the Pseudo-Basis: a Warm-Up

We present here a new method of communicating the pseudo-basis, which is a straightforward implementation of the generalized broadcasting technique.

The key point is the following observation:

Lemma 4. *Let $\mathcal{W} = (\mathbf{y}^{(i)} : i \in I)$ be a pseudo-basis of the set of received vectors; then if Bob knows m elements of \mathcal{W} , he knows at least m channels that have been forged by Eve.*

Proof. By subtracting the original codeword from an element of the pseudo-basis, Bob knows the corresponding error; furthermore, these errors form a basis of the entire error space (Remark 1). Now if Bob knows m elements of the pseudo-basis, he then knows m of these errors, which necessarily affect at least m coordinates since they are linearly independent. The claim then follows. \square

The sub-protocol consisting of the transmission of the pseudo-basis by Alice is simply the following:

Protocol 2. Alice wishes to communicate to Bob a pseudo-basis \mathcal{W} of cardinality w .

For any $i = 1, \dots, w$, she then uses $(i - 1)$ -generalized broadcast to communicate the i -th element of the pseudo-basis to Bob.

Lemmas 3 and 4 ensure that this technique is secure; we now compute its cost:

- Each element of the pseudo-basis is a vector of \mathbb{F}_q^n ;
- using m -generalized broadcast to communicate n elements of \mathbb{F}_q requires communicating $\left\lceil \frac{n}{m+1} \right\rceil n$ field elements;
- hence Protocol 2 requires communicating the following number of elements of \mathbb{F}_q :

$$\sum_{i=1}^w \left\lceil \frac{n}{i} \right\rceil n = O \left(n^2 \sum_{i=1}^w \frac{1}{i} \right) = O(n^2 \log w)$$

which means that we have reduced to $O(n^2 \log w)$ the total communication complexity.

This complexity is still one logarithmic factor short of our goal; in the next section we show a more advanced technique that allows to bring down the cost to $O(n^2)$ field elements.

5.3 Improved Transmission of the Pseudo-Basis: the Final Version

In this section we show a more advanced technique to communicate the pseudo-basis. The key idea is the following: denote by w the size of the pseudo-basis; if Alice can find a received word \mathbf{y} which is subject to an error of weight cw for some constant c and sends it to Bob, then Bob will learn the position of at least cw corrupted channels. Alice will thus be able to use cw -generalized broadcast as in Lemma 3 to communicate the elements of the pseudo-basis (which amount to wn symbols); since cw -generalized broadcast of a symbol has a cost of $O(n/cw)$, the total cost of communicating the pseudo-basis will thus be $(wn) \cdot O(n/cw) = O(n^2)$.

We thus devise an algorithm that allows Alice to find a word \mathbf{y} subject to at least $m = \Omega(w)$ errors (for instance, such condition is met if \mathbf{y} is subject to $\Omega(t)$ errors, since $w \leq t$). Notice that

such a word \mathbf{y} may not exist among the received words $\{\mathbf{y}^{(i)}\}$, therefore we will look for a linear combination of the $\mathbf{y}^{(i)}$ with this property.

As mentioned in Sections 2 and 3, Alice will make extensive use of a decoding algorithm. Recall that a code of distance d can be uniquely decoded from up to $\lfloor (d-1)/2 \rfloor$ errors, and that in the case of Reed-Solomon codes, such decoding can be performed in time polynomial in n [13]; this means that for any Reed-Solomon code \mathcal{C} there exists an algorithm that takes as input a word $\mathbf{y} \in \mathbb{F}_q^n$ and outputs a decomposition $\mathbf{y} = \mathbf{x} + \mathbf{e}$ with $\mathbf{x} \in \mathcal{C}$ and $w_H(\mathbf{e}) \leq \lfloor (d-1)/2 \rfloor$ (if such a decomposition does not exist, the algorithm outputs an error message \perp).

Protocol 3. Alice has received the words $\mathbf{y}^{(1)}, \dots, \mathbf{y}^{(r)}$ and has computed a pseudo-basis $\{\mathbf{y}^{(i)} : i \in I\}$ of them; denote by w its cardinality. Alice proceeds with the following actions:

- she uses Algorithm 1 below to find a “special word” \mathbf{y} , with coefficients $(\mu_i : i \in I)$ such that $\mathbf{y} = \sum_{i \in I} \mu_i \mathbf{y}^{(i)}$. She then communicates to Bob the triplet $(I, (\mu_i : i \in I), \mathbf{y})$ by using ordinary broadcast.
- Finally, she communicates the pseudo-basis of the received values by using m -generalized broadcast, where $m := \min(w, t/3)$, w being the cardinality of the pseudo-basis.

Before describing the algorithm formally and proving its validity, we sketch the idea. Alice has computed a pseudo-basis $\{\mathbf{y}^{(i)} : i \in I\}$. For $i \in I$, she applies the decoding algorithm to $\mathbf{y}^{(i)} = \mathbf{x}^{(i)} + \mathbf{e}^{(i)}$. If the decoding algorithm fails, it means that $\mathbf{y}^{(i)}$ is at a large Hamming distance from any codeword, in particular from Bob’s codeword $\mathbf{x}^{(i)}$, and the single $\mathbf{y}^{(i)}$ is the required linear combination. If the decoding algorithm succeeds for every i , Alice obtains decompositions

$$\mathbf{y}^{(i)} = \tilde{\mathbf{x}}^{(i)} + \tilde{\mathbf{e}}^{(i)}$$

where $\tilde{\mathbf{x}}^{(i)}$ is some codeword. Alice must be careful, because she has no guarantee that the codeword $\tilde{\mathbf{x}}^{(i)}$ coincides with Bob’s codeword $\mathbf{x}^{(i)}$, and hence that $\tilde{\mathbf{e}}^{(i)}$ coincides with Eve’s error vector $\mathbf{e}^{(i)}$. What Alice then does is look for a linear combination $\sum_i \mu_i \tilde{\mathbf{e}}^{(i)}$ that has Hamming weight at least $t/3$ and at most $2t/3$. If she is able to find one, then a simple Hamming distance argument guarantees that the corresponding linear combination of Eve’s original errors $\sum_i \mu_i \mathbf{e}^{(i)}$ also has Hamming weight at least $t/3$. If Alice is unable to find such a linear combination, then she falls back on constructing one that has weight not more than $2t/3$ and at least the cardinality w of the pseudo-basis. This will yield an alternative form of the desired result. We now describe this formally.

Algorithm 1. Alice has a pseudo-basis $(\mathbf{y}^{(i)} : i = 1, \dots, w)$ (indices have been changed to simplify the notation); the algorithm allows Alice to identify a word \mathbf{y} subject to at least $m := \min(w, t/3)$ errors introduced by Eve.

In the following steps, whenever we say that the output of the algorithm is a word $\mathbf{y}^{(i)}$, we implicitly assume that the algorithm also outputs the index i ; more generally, whenever the

algorithm outputs a linear combination $\sum_i \mu_i \mathbf{y}^{(i)}$ of the words in the pseudo-basis, we assume that it also outputs the coefficient vector (μ_1, \dots, μ_w) of the linear combination.

1. Alice uses a unique decoding algorithm to decode the elements of the pseudo-basis; if the algorithm fails for a given word $\mathbf{y}^{(i)}$ (i.e., it doesn't output a codeword having distance at most $t/2$ from $\mathbf{y}^{(i)}$), then Algorithm 1 stops and outputs $\mathbf{y}^{(i)}$.
2. If the decoding algorithm worked for every i , Alice gets a decomposition $\mathbf{y}^{(i)} = \tilde{\mathbf{x}}^{(i)} + \tilde{\mathbf{e}}^{(i)}$ with $\tilde{\mathbf{x}}^{(i)} \in \mathcal{C}$ and $w_H(\tilde{\mathbf{e}}^{(i)}) \leq t/2$ for every i ; notice that it is not guaranteed that the $\tilde{\mathbf{x}}^{(i)}$ coincide with the codewords $\mathbf{x}^{(i)}$ originally chosen by Bob.
If any of the $\tilde{\mathbf{e}}^{(i)}$ has weight greater than $t/3$, the algorithm stops and outputs $\mathbf{y}^{(i)}$.
3. Define $\tilde{\mathbf{f}}^{(1)} := \tilde{\mathbf{e}}^{(1)}$ and $\tilde{\mathbf{y}}^{(1)} := \mathbf{y}^{(1)}$. For any $i = 2, \dots, w$, proceed with the following actions:

- let $\lambda^{(i)}$ be a non-zero element of \mathbb{F}_q such that $\tilde{\mathbf{f}}_j^{(i-1)} + \lambda^{(i)} \tilde{\mathbf{e}}_j^{(i)} \neq 0$ for any coordinate $j \in \{1, 2, \dots, n\}$ for which $\tilde{\mathbf{f}}_j^{(i-1)} \neq 0$.
- let $\tilde{\mathbf{f}}^{(i)} := \tilde{\mathbf{f}}^{(i-1)} + \lambda^{(i)} \tilde{\mathbf{e}}^{(i)}$ and $\tilde{\mathbf{y}}^{(i)} := \tilde{\mathbf{y}}^{(i-1)} + \lambda^{(i)} \mathbf{y}^{(i)}$;
if $w_H(\tilde{\mathbf{f}}^{(i)}) > t/3$, stop and output $\tilde{\mathbf{y}}^{(i)}$.

4. Output $\tilde{\mathbf{y}}^{(w)}$.

We can now prove that this algorithm allows Alice to find the desired codeword, which naturally implies that Protocol 3 indeed allows for reliable communication of the pseudo-basis:

Proposition 3. *Algorithm 1 allows Alice to find a word \mathbf{y} subject to an error introduced by Eve of weight at least $m := \min(w, t/3)$.*

Proof. The following observation is the key point of the algorithm:

Lemma 5. *Let $\mathbf{y} = \mathbf{x} + \mathbf{e} = \tilde{\mathbf{x}} + \tilde{\mathbf{e}}$ for $\mathbf{x}, \tilde{\mathbf{x}} \in \mathcal{C}$. Then if $\tilde{\mathbf{e}}$ satisfies $w_H(\tilde{\mathbf{e}}) \leq 2t/3$, we have that $w_H(\mathbf{e}) \geq \min\{w_H(\tilde{\mathbf{e}}), t/3\}$.*

Proof. The claim is trivial if $\mathbf{e} = \tilde{\mathbf{e}}$; hence assume that $\mathbf{e} \neq \tilde{\mathbf{e}}$. Notice that $\mathbf{e} - \tilde{\mathbf{e}} = \tilde{\mathbf{x}} - \mathbf{x}$; hence since $d_{\min}(\mathcal{C}) = t + 1$, we have that

$$t + 1 \leq w_H(\mathbf{e} - \tilde{\mathbf{e}}) \leq w_H(\mathbf{e}) + w_H(\tilde{\mathbf{e}}) \leq w_H(\mathbf{e}) + \frac{2t}{3}$$

Hence we have that $w_H(\mathbf{e}) \geq t/3$, so that the claim is proved. □

We now analyze the algorithm step-by-step:

1. if decoding fails for a word $\mathbf{y}^{(i)}$, then it is guaranteed that the error introduced by Eve on it has weight bigger than $t/2 > m$ (otherwise, the unique decoding algorithm would succeed since $d_{\min}(\mathcal{C}) = t + 1$).
2. since by assumption $w_{\mathbf{H}}(\tilde{\mathbf{e}}^{(i)}) \leq t/2 \leq 2t/3$, if we also have $t/3 \leq w_{\mathbf{H}}(\tilde{\mathbf{e}}^{(i)})$, then thanks to Lemma 5 the output $\mathbf{y}^{(i)}$ is of the desired type.
3. Since the algorithm did not abort at step 2, all elements $\tilde{\mathbf{e}}^{(i)}$ have weight at most $t/3$.

First notice that if the algorithm did not produce $\tilde{\mathbf{f}}^{(i-1)}$ as output, then $\tilde{\mathbf{f}}^{(i)}$ is well-defined: indeed, we have that $w_{\mathbf{H}}(\tilde{\mathbf{f}}^{(i-1)}) \leq t/3$; this means that $\lambda^{(i)}$ is well-defined, since it is an element of \mathbb{F}_q that has to be different from 0 and from at most $t/3 < n - 1$ elements.

Now if the algorithm outputs $\tilde{\mathbf{f}}^{(i)}$, then necessarily $w_{\mathbf{H}}(\tilde{\mathbf{f}}^{(i-1)}) \leq t/3$ (otherwise the algorithm would have stopped before computing $\tilde{\mathbf{f}}^{(i)}$); furthermore, by assumption we have that $w_{\mathbf{H}}(\tilde{\mathbf{e}}^{(i)}) \leq t/3$, so that $w_{\mathbf{H}}(\tilde{\mathbf{f}}^{(i)}) \leq 2t/3$ and we can apply Lemma 5, so that the output is of the desired type.

4. Notice that for any $i = 1, \dots, w$, we have that $\tilde{\mathbf{f}}^{(i)}$ has maximal weight among elements of the vector space $\langle \tilde{\mathbf{e}}^{(1)}, \dots, \tilde{\mathbf{e}}^{(i)} \rangle$ (the condition on $\lambda^{(i)}$ ensures that this condition is met at each step). Hence since the elements $\{\tilde{\mathbf{e}}^{(1)}, \dots, \tilde{\mathbf{e}}^{(w)}\}$ are linearly independent (because their syndromes are linearly independent, since $(\mathbf{y}^{(1)}, \dots, \mathbf{y}^{(w)})$ is a pseudo-basis), we have that $w_{\mathbf{H}}(\tilde{\mathbf{f}}^{(i)}) \geq i$ for any i .

In particular, we have that $w_{\mathbf{H}}(\tilde{\mathbf{f}}^{(w)}) \geq w$; hence since $w_{\mathbf{H}}(\tilde{\mathbf{f}}^{(w)}) \leq 2t/3$ as remarked above, we have that the output $\tilde{\mathbf{y}}^{(w)}$ is of the desired type.

□

Remark 2. Protocol 3 requires Alice to use ordinary broadcast to communicate a single vector of \mathbb{F}_q^n (hence transmitting n^2 elements of \mathbb{F}_q), then to use m -generalized broadcast with $m \geq \min\{w, t/3\}$ to communicate $w \leq t$ vectors of \mathbb{F}_q^n (hence transmitting at most $3n^2$ elements of \mathbb{F}_q). We thus get a total of at most $4n^2$ elements of \mathbb{F}_q to be transmitted.

Furthermore, Algorithm 1 has running time polynomial in n , as long as the code \mathcal{C} has a unique-decoding algorithm of polynomial running time as well. As already remarked, such algorithms exist for instance for Reed-Solomon codes.

We study the second bottleneck of the original protocol in the next section.

5.4 The Improved Communication of the Masked Secrets

We present in this section the second key improvement to the protocol: after the pseudo-basis is communicated, we devise a way to lower the cost of transmitting to Bob the masked secrets and the information to open the masks. We aim at a cost linear in the number ℓ of secrets to be transmitted (while it was quadratic in Protocol 1). As in Section 5.3, Alice makes use of a unique decoding algorithm.

Protocol 4. The protocol is performed once the pseudo-basis has been communicated to Bob; we thus assume that Bob knows the global support $\mathcal{S} := \cup_i \text{support}(\mathbf{e}^{(i)})$ of the errors affecting the elements $\mathbf{y}^{(i)}$ (cf. Remark 1). We assume that Alice wishes to communicate ℓ secret elements $s^{(1)}, \dots, s^{(\ell)}$ of \mathbb{F}_q to Bob, and that ℓ codewords $\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(\ell)}$ of \mathcal{C} have been sent by Bob to Alice (who has received $\mathbf{y}^{(1)}, \dots, \mathbf{y}^{(\ell)}$) and have not been disclosed in other phases.

- Alice uses a unique decoding algorithm to decode $\mathbf{y}^{(i)}$, so that for every i she obtains (if decoding was successful) a decomposition $\mathbf{y}^{(i)} = \tilde{\mathbf{x}}^{(i)} + \tilde{\mathbf{e}}^{(i)}$ with $\tilde{\mathbf{x}}^{(i)} \in \mathcal{C}$ and $w_{\mathbf{H}}(\tilde{\mathbf{e}}^{(i)}) \leq t/2$.

For every $i = 1, \dots, \ell$ she then communicates the following elements to Bob:

- the syndrome $\mathbf{H}(\mathbf{y}^{(i)})^T$ via $t/2$ -generalized broadcast;
- the elements $z_1^{(i)}, z_2^{(i)}$ of \mathbb{F}_q by ordinary broadcast, where

$$\begin{aligned} z_1^{(i)} &:= s^{(i)} + \mathbf{h}(\mathbf{y}^{(i)})^T \\ z_2^{(i)} &:= \begin{cases} s^{(i)} + \mathbf{h}(\tilde{\mathbf{x}}^{(i)})^T & \text{if decoding succeeded,} \\ 0 & \text{otherwise.} \end{cases} \end{aligned}$$

- Bob can then obtain each secret $s^{(i)}$ in a different way depending on the size of the global support \mathcal{S} of the errors:
 - if $|\mathcal{S}| \geq t/2$, he uses the knowledge of the syndrome of $\mathbf{y}^{(i)}$ and of the support of the error to compute $\mathbf{y}^{(i)}$, so that he can compute $z_1^{(i)} - \mathbf{h}(\mathbf{y}^{(i)})^T$ as well.
 - if $|\mathcal{S}| < t/2$, he ignores the syndrome that has been communicated to him, and computes $z_2^{(i)} - \mathbf{h}(\mathbf{x}^{(i)})^T$.

We now prove that this protocol works and is secure:

Proposition 4. *The above protocol allows for private and reliable communication of ℓ elements of \mathbb{F}_q .*

Proof. We check Privacy and Reliability.

Privacy: we have already observed in Proposition 2 that Eve has no information on $\mathbf{h}\mathbf{y}^T$ (we drop the index (i) to simplify notation), so that z_1 perfectly hides the secret. Now notice that if \mathbf{y} can be decoded, then $z_2 = s + \mathbf{h}\tilde{\mathbf{x}}^T = z_1 - \mathbf{h}\tilde{\mathbf{e}}^T$; hence to conclude, it suffices to prove that Eve already knows whether \mathbf{y} can be decoded or not, and that she knows $\tilde{\mathbf{e}}$ if \mathbf{y} can be decoded. We prove this claim in the following lemma:

Lemma 6. *Let \mathbf{x} be a codeword sent by Bob to Alice, and let $\mathbf{y} = \mathbf{x} + \mathbf{e}$ be the received vector. Then Eve knows whether \mathbf{y} can be decoded (i.e. $\mathbf{y} = \tilde{\mathbf{x}} + \tilde{\mathbf{e}}$ as above) or not; furthermore, if \mathbf{y} can be decoded, then she knows $\tilde{\mathbf{e}}$.*

Proof. By definition, $\tilde{\mathbf{e}}$ is a vector of minimum weight (and of weight at most $t/2$) such that $\mathbf{y} - \tilde{\mathbf{e}}$ belongs to \mathcal{C} ; notice that the last condition is equivalent to require that $\mathbf{e} - \tilde{\mathbf{e}}$ belongs to \mathcal{C} . Now these requirements uniquely determine $\tilde{\mathbf{e}}$: indeed, if by contradiction $\mathbf{e} - \mathbf{e}' \in \mathcal{C}$ for another \mathbf{e}' , then $\mathbf{e}' - \tilde{\mathbf{e}}$ would belong to \mathcal{C} , a contradiction since $w_{\mathbb{H}}(\mathbf{e}' - \tilde{\mathbf{e}}) \leq t/2 + t/2 < d_{\min}(\mathcal{C})$.

Hence $\tilde{\mathbf{e}}$ is uniquely determined by \mathbf{e} and \mathcal{C} : Eve can thus compute it from the data in her possession. Notice that, in particular, she knows whether $\tilde{\mathbf{e}}$ exists or not, i.e. whether decoding of \mathbf{y} is possible or not. □

Reliability: we have two possible cases:

- if $|\mathcal{S}| \geq t/2$, then Bob is able to acquire the syndrome $\mathbf{H}\mathbf{y}^T$ of \mathbf{y} via $t/2$ -generalized broadcast (cf. Lemma 3); thus as remarked in Proposition 2, he can recover \mathbf{y} and open the mask to get the secret.
- if $|\mathcal{S}| < t/2$, then Bob knows that Alice has correctly decoded \mathbf{y} , since Eve introduced less than $d_{\min}/2$ errors; thus $\tilde{\mathbf{x}} = \mathbf{x}$ so that $z_2 - \mathbf{h}\mathbf{x}^T = (s + \mathbf{h}\tilde{\mathbf{x}}^T) - \mathbf{h}\mathbf{x}^T = s$.

Notice that in this case Bob will have failed to decode the $t/2$ -generalized broadcast but he will simply ignore the elements received in this way. □

Remark 3. Notice that we could further improve the efficiency of this protocol by requiring Alice to use w -generalized broadcast (instead of regular one) to communicate the elements $z_1^{(i)}$ and $z_2^{(i)}$, where w is the size of the pseudo-basis; this, however, would not reduce the order of magnitude of the total cost.

6 The Improved Protocol

The improved protocol simply implements the new techniques of sections 5.3 and 5.4.

Protocol 5. The protocol allows Alice to communicate ℓ secret elements $s^{(1)}, \dots, s^{(\ell)}$ of \mathbb{F}_q to Bob, where q is an arbitrary integer with $q > n$. The protocol takes as input an MDS code \mathcal{C} of parameters $[n, t+1, t+1]_q$ and a vector \mathbf{h} of length n as in Lemma 1.

- I. Bob chooses $t + \ell + 1$ uniformly random and independent codewords $\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(t+\ell+1)}$ of \mathcal{C} and sends them over the channels to Alice.
- II. Alice receives the corrupted versions $\mathbf{y}^{(1)}, \dots, \mathbf{y}^{(t+\ell+1)}$, and she computes a pseudo-basis $\{\mathbf{y}^{(i)} : i \in I\}$ of the received values; she then proceeds with the following actions:
 - (i) She uses Protocol 3 to communicate the pseudo-basis to Bob.
 - (ii) She then uses the remaining words to communicate to Bob the masked secrets and the data to retrieve them as in the first part of Protocol 4.

III. Upon receiving the pseudo-basis, Bob proceeds to compute the global support \mathcal{S} of the error space; he can then obtain each secret $s^{(i)}$ as specified in the corresponding part of Protocol 4.

Notice that privacy and reliability of the protocol follow from the previous discussions; we now analyze the complexity of the protocol:

Communication complexity: we can set $\ell := 1$.

- Step I requires transmitting $t + 2$ words of \mathbb{F}_q^n over the channels, thus requiring a total of $O(n^2)$ symbols to be transmitted.
- Step II-(i) requires transmitting $O(n^2)$ elements of \mathbb{F}_q as shown in Remark 2.
- Finally, step II-(ii) requires using $t/2$ -generalized broadcast to communicate n symbols, and standard broadcast to communicate 2 symbols, thus giving a total of $O(n)$ elements to be transmitted.

Hence since we can assume that $q = O(n)$, we get a total communication complexity of

$$O(n^2 \log n)$$

bits to be transmitted to communicate a single-bit secret.

Transfer rate: optimal rate is achieved for $\ell = \Omega(n)$.

- Step I requires transmitting $t + \ell + 1 = \ell + O(n)$ codewords, for a total of $n\ell + O(n^2)$ symbols.
- Step II-(i) remains unchanged from the single-bit case, and thus requires transmitting $O(n^2)$ symbols.
- Finally, step II-(ii) uses $t/2$ -generalized broadcast to communicate ℓt elements of \mathbb{F}_q and standard broadcast to communicate 2ℓ elements of \mathbb{F}_q , so that the overall cost is equal to $4n\ell$ symbols to be transmitted.

To sum up, the overall transmission rate is equal to

$$\frac{5n\ell + O(n^2)}{\ell} = 5n + O(n^2/\ell)$$

Furthermore, by using Reed-Solomon codes (instead of arbitrary MDS ones), we then have that Protocol 5 has computational cost polynomial in n for both Alice and Bob.

7 Generalization to Linear Combinations of Errors and Eavesdropped Data

We present in this section a more general scenario, where Alice and Bob can communicate words to each other in the presence of a more powerful adversary Eve, who can eavesdrop linear combinations of the transmitted symbols and inject linear combinations of errors from a set of her choice. We show that Protocol 1 (our vanilla protocol) can be generalized in order to provide security in this more complex scenario; this shows how our protocol carries greater potential for more general settings compared to previous work that relied on more cumbersome communication techniques.

The generalization of the communication model is defined as follows:

Definition 2 (Generalized Model). Let m, n, t be integers, and \mathbb{F}_q be a finite field. Alice and Bob can communicate to each other n -tuples $\mathbf{x} = [\mathbf{x}_1, \dots, \mathbf{x}_n]$ where each \mathbf{x}_i is a column vector in \mathbb{F}_q^m .

The adversary Eve is computationally unbounded, and selects at the beginning of the protocol t “eavesdropping vectors” $\boldsymbol{\lambda}^{(1)}, \dots, \boldsymbol{\lambda}^{(t)} \in \mathbb{F}_q^m$ and t “tampering vectors” $\boldsymbol{\mu}^{(1)}, \dots, \boldsymbol{\mu}^{(t)} \in \mathbb{F}_q^n$. Whenever $\mathbf{x} = [\mathbf{x}_1, \dots, \mathbf{x}_n]$ is transmitted (either from Alice to Bob or from Bob to Alice), the following happens:

- Eve learns the value of

$$\boldsymbol{\lambda}^{(i)} \mathbf{x}^T = \lambda_1^{(i)} \mathbf{x}_1 + \dots + \lambda_n^{(i)} \mathbf{x}_n$$

for every $i = 1, \dots, t$.

- Eve selects t columns vectors $\boldsymbol{\Delta}^{(1)}, \dots, \boldsymbol{\Delta}^{(t)} \in \mathbb{F}_q^m$; the intended receiver gets the message $\mathbf{y} = \mathbf{x} + \mathbf{e}$, where

$$\mathbf{e} = \sum_{i=1, \dots, t} \boldsymbol{\Delta}^{(i)} \otimes \boldsymbol{\mu}^{(i)} = \sum_{i=1, \dots, t} \left[\Delta^{(i)} \mu_1^{(i)}, \dots, \Delta^{(i)} \mu_n^{(i)} \right].$$

Notice that PSMT can be seen as a more restrictive version of this model, where Eve is forced to choose vectors $\boldsymbol{\lambda}$ and $\boldsymbol{\mu}$ of weight 1, and where moreover $\boldsymbol{\lambda}^{(i)} = \boldsymbol{\mu}^{(i)}$ for every i .

Though the adversary of Definition 2 is way more powerful than the one in classical PSMT, we show that Protocol 1 can be adapted to provide security in this scenario as well. More precisely, the only relevant modification we have to perform is the following: instead of classical codes, which can correct from errors with bounded Hamming weight, we use *rank* codes which can correct from errors with bounded rank; notice that such are the errors introduced by the adversary in Definition 2.

As a final remark, we stress the fact that this more complex scenario is not a purely gratuitous generalization but has an application to Secure Network Coding.

Indeed, in Network Coding, one or several transmitters are connected to one or several receivers by a network, i.e. a directed multigraph with source nodes and destination nodes; “network coding” means that each node performs \mathbb{F}_q -linear operations on the symbols received via the incoming edges, and sends the results through the outbound edges. Each sender can feed input to its corresponding source node, and each receiver can read the output of the corresponding destination node. In

the “Secure Network Coding” scenario (e.g. [18]), we have typically a single transmitter and an adversary that controls t edges, meaning he can read the symbols transmitted over these edges and replace them by symbols of his choice.

In case there is a single receiver, we can identify the sender with Alice and the receiver with Bob; it is then readily seen that the communication between Alice and Bob is affected by the adversary precisely as in Definition 2. This means that our protocol can be used to provide security for a Network Coding scenario with a single receiver, as long as communication is also possible from the receiver to the sender. Though Network Coding was originally introduced in a multicast scenario [2, 11], it has since been proved useful in single sender - single receiver scenarios as well [12, 21].

Until very recently, existing work on Secure Network Coding assumes that information can only flow from sender to receiver. Notably, the work of Silva and Kschischang [18] presents a one-round protocol that is secure as long as $t < n/3$; recently, the present authors introduced a protocol [19] that achieves security for any $t < n/2$, in a multiple receiver context, by allowing communication from receiver to sender as well. This protocol uses three rounds of communication. We sketch below how a two-round generalization of Protocol 1 is also secure for any $t < n/2$. This can be directly applied to Secure Network Coding in a unicast (single transmitter – single receiver) scenario.

We make the final remark that Jaggi et al. [7] studied a similar case, where the adversary controls vertices instead of edges of the network; their protocol lets the adversary inject up to $n/2$ errors, but with a weaker notion of security in that it must drop the privacy requirement and the reconstruction process admits a positive error probability.

7.1 Communication with Rank-Metric Codes

We show in this section how the machinery of Protocol 1 can be adapted to another type of code, defined under the *rank metric* [3], [5]. These have been extensively used in Secure Network Coding [17], [18].

Definition 3 (Rank-Metric Code). Given the \mathbb{F}_q -linear space $\mathbb{F}_q^{m \times n}$ (m -by- n matrices over \mathbb{F}_q), we can define the *rank distance* between its elements by letting $d_R(\mathbf{x}, \mathbf{y}) := \text{rank}(\mathbf{y} - \mathbf{x})$.

A *rank-metric* code \mathcal{C} is a non-empty subset of $\mathbb{F}_q^{m \times n}$ with induced rank distance; by identifying the field \mathbb{F}_{q^m} with \mathbb{F}_q^m , we can view \mathcal{C} as a code over \mathbb{F}_{q^m} , and require it to be *linear* over \mathbb{F}_{q^m} ; we can hence speak of *block length* and *dimension* of such a code (as in the Hamming case) and of *minimum (rank) distance*, and combine these parameters into the triplet $[n, k, d]_{q^m}$.

The equivalent concept of MDS in this setting is called *Maximum Rank-Distance*: a rank-metric code is Maximum Rank-Distance (or *MRD* for short) if $k + d = n + 1$; an MRD code of arbitrary dimension k and length n exists if and only if $m \geq n$ [3]. More precisely, for any q and any m, n, k with $m \geq n$, we can construct a *Gabidulin code* [5] of parameters $[n, k, n - k + 1]_{q^m}$.

As in the Hamming case, we can express a rank-metric code \mathcal{C} in term of a *parity-check* matrix \mathbf{H} , i.e.

$$\mathcal{C} = \{\mathbf{x} \in \mathbb{F}_{q^m}^n : \mathbf{H}\mathbf{x}^T = \mathbf{0}\}.$$

Furthermore, we can define the associated *syndrome* map $\sigma : \mathbf{w} \mapsto \mathbf{H}\mathbf{w}^T$; clearly, we have that $\mathcal{C} = \ker(\sigma)$.

As a first building block, we construct a rank-metric equivalent of Lemma 1: we show how to define a rank-metric code \mathcal{C} and a vector \mathbf{h} such that the adversary has no information of the value of $\mathbf{h}\mathbf{x}^T$ if a random $\mathbf{x} \in \mathcal{C}$ is transmitted between Alice and Bob. Here the key point is that the vector \mathbf{h} has entries in \mathbb{F}_{q^m} rather than in \mathbb{F}_q .

Lemma 7. *For any q , n , $m \geq n + 1$ and $t < n$ there exists an MRD code \mathcal{C} of parameters $[n, t + 1, n - t]_{q^m}$ and a vector $\mathbf{h} \in \mathbb{F}_{q^m}^n$ such that given a random codeword $\mathbf{x} \in \mathcal{C}$ and t arbitrary vectors $\boldsymbol{\lambda}^{(1)}, \dots, \boldsymbol{\lambda}^{(t)} \in \mathbb{F}_q^n$, the value of $\mathbf{h}\mathbf{x}^T$ is completely undetermined even if $\boldsymbol{\lambda}^{(1)}\mathbf{x}^T, \dots, \boldsymbol{\lambda}^{(t)}\mathbf{x}^T$ are known.*

Proof. We construct \mathcal{C} and \mathbf{h} by adapting the blueprint of Lemma 1 to the rank-metric setting: hence we first let \mathcal{C}' be an MRD code of parameters $[n + 1, t + 1, n - t + 1]$, and let

$$\mathcal{C} := \{\mathbf{x} \in \mathbb{F}_{q^m}^n : \exists x_{n+1} \in \mathbb{F}_{q^m} \text{ with } (\mathbf{x}, x_{n+1}) \in \mathcal{C}'\}$$

We then define \mathbf{h} by selecting a parity-check matrix \mathbf{H}' of \mathcal{C}' , and choosing a row thereof of the form

$$(\mathbf{h}, \alpha) \in \mathbb{F}_{q^m}^{n+1} \text{ with } \mathbf{h} \in \mathbb{F}_{q^m}^n, 0 \neq \alpha \in \mathbb{F}_{q^m}.$$

Now consider the following matrix:

$$\mathbf{M} := \left[\begin{array}{c|c} \mathbf{H}' & \\ \hline \boldsymbol{\lambda}^{(1)} & 0 \\ \vdots & \vdots \\ \boldsymbol{\lambda}^{(t)} & 0 \\ \hline \mathbf{0} & -\alpha \end{array} \right] \cdot \begin{bmatrix} \mathbf{x}_1 \\ \vdots \\ \mathbf{x}_n \\ x_{n+1} \end{bmatrix}$$

clearly, if M is non-singular then the claim holds (we assume as a worst-case scenario that the vectors $(\boldsymbol{\lambda}^{(i)} : i = 1, \dots, t)$ are linearly independent). Now by properties of MRD codes, the matrix $\begin{bmatrix} \mathbf{H} \\ \mathbf{B} \end{bmatrix}$ is non-singular for any full-rank matrix $\mathbf{B} \in \mathbb{F}_q^{t \times n}$ [17]; hence in particular \mathbf{M} is of full-rank, so that the claim holds. □

In a symmetric fashion, we now present a rank-metric version of the broadcast protocol, hence allowing for reliable communication of messages, although with no guarantee of privacy.

Lemma 8. *Given any q and $n \leq m$, let $\mathcal{C}_{\text{BCAST}}$ be an MRD code of parameters $[n, 1, n]_{q^m}$. Then if an arbitrary $\mathbf{x} \in \mathcal{C}_{\text{BCAST}}$ is transmitted between the players, the receiver can always recover \mathbf{x} from the received message \mathbf{y} by computing the closest codeword to \mathbf{y} .*

Proof. By Definition 2, the receiver obtains a message \mathbf{y} with $\mathbf{y} = \mathbf{x} + \mathbf{e}$, where the error \mathbf{e} introduced by Eve is of rank at most t . Hence the original codeword \mathbf{x} can be recovered since $\mathcal{C}_{\text{BCAST}}$ has rank-distance $n \geq 2t + 1$. □

Furthermore, we see that the machinery of the pseudo-basis can be adapted to the rank-metric case; we begin with the equivalent of Lemma 2:

Lemma 9. *Let \mathcal{C} be a rank-metric code of parameters $[n, k, d]_{q^m}$; let \mathbf{H} be a parity-check matrix of \mathcal{C} , and let $W \leq \mathbb{F}_{q^m}^n$ be an \mathbb{F}_{q^m} -linear subspace with the property that each $\mathbf{w} \in W$ is of rank at most $d - 1$ over \mathbb{F}_q . Then the syndrome map $\sigma : \mathbf{w} \mapsto \mathbf{H}\mathbf{w}^T$ is injective on W .*

As in the Hamming case, given a rank-metric code \mathcal{C} of length n over \mathbb{F}_q and a set \mathcal{Y} of vectors in $\mathbb{F}_{q^m}^n$, we call a pseudo-basis of \mathcal{Y} a subset $\mathcal{W} \subseteq \mathcal{Y}$ such that $\sigma(\mathcal{W})$ is a basis of $\langle \sigma(\mathcal{Y}) \rangle$, where σ denotes a syndrome map of \mathcal{C} .

Again, we have that since the codomain of σ is equal to $\mathbb{F}_{q^m}^{n-k}$, a pseudo-basis has cardinality at most $n - k$; furthermore, a pseudo-basis can be computed in time polynomial in n .

We now show that if the set \mathcal{Y} consists of corrupted codewords, affected by errors introduced by the adversary, then a pseudo-basis corresponds to a basis of the error space in this setting as well; the proof follows the same steps as that of Proposition 1:

Proposition 5. *Let \mathcal{C} be a linear rank-metric code of parameters $[n, k, d]_{q^m}$, and let $\mathcal{X}, \mathcal{E}, \mathcal{Y}$ be three subsets:*

$$\begin{aligned} \mathcal{X} &:= \{\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(r)}\} \subseteq \mathcal{C}, \\ \mathcal{E} &:= \{\mathbf{e}^{(1)}, \dots, \mathbf{e}^{(r)}\} \subseteq \mathbb{F}_q^n \quad \text{such that } \text{rank}(\mathbf{e}) \leq d - 1 \text{ for all } \mathbf{e} \in \langle \mathcal{E} \rangle_{\mathbb{F}_{q^m}}, \\ \mathcal{Y} &:= \{\mathbf{y}^{(1)}, \dots, \mathbf{y}^{(r)}\} \subseteq \mathbb{F}_q^n \quad \text{with } \mathbf{y}^{(j)} = \mathbf{x}^{(j)} + \mathbf{e}^{(j)} \text{ for every } j \end{aligned}$$

Then, given knowledge of \mathcal{X} and a pseudo-basis of \mathcal{Y} , we can compute $\mathbf{e}^{(j)}$ from its syndrome $\sigma(\mathbf{e}^{(j)})$, for any $1 \leq j \leq r$.

We show in the following section how to implement the techniques we presented.

7.2 The Protocol for The Rank-Metric Case

We define in this section the protocol for private and reliable communication in the setting of Definition 2.

Protocol 6. We assume that Alice and Bob can communicate vectors to each other in the presence of an adversary Eve as in Definition 2; we assume that $t < n/2$ and that $m > n$.

The protocol allows Alice to communicate ℓ secret elements $s^{(1)}, \dots, s^{(\ell)}$ of \mathbb{F}_{q^m} to Bob, and it takes as input a pair $(\mathcal{C}, \mathbf{h})$ as in Lemma 7 for private communication, and a rank-metric code $\mathcal{C}_{\text{BCAST}}$ for reliable communication as in Lemma 8.

- I. Bob chooses $t + \ell$ uniformly random and independent codewords $\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(t+\ell)}$ of \mathcal{C} and communicates them to Alice.
- II. Alice receives the corrupted versions $\mathbf{y}^{(1)} = \mathbf{x}^{(1)} + \mathbf{e}^{(1)}, \dots, \mathbf{y}^{(t+\ell)} = \mathbf{x}^{(t+\ell)} + \mathbf{e}^{(t+\ell)}$; she then proceeds with the following actions:
 - (i) She computes a pseudo-basis $\{\mathbf{y}^{(i)} : i \in I\}$ for $I \subset \{1, \dots, t + \ell\}$ of the received values, and uses $\mathcal{C}_{\text{BCAST}}$ to reliably communicate to Bob $(i, \mathbf{y}^{(i)} : i \in I)$.

(ii) She then considers the first ℓ words that do *not* belong to the pseudo-basis; to ease the notation, we will re-name them $\mathbf{y}^{(1)}, \dots, \mathbf{y}^{(\ell)}$. For each secret $s^{(j)}$ to be communicated she broadcasts to Bob the following two elements:

- $\mathbf{H}(\mathbf{y}^{(j)})^T$, the syndrome of $\mathbf{y}^{(j)}$;
- $s^{(j)} + \mathbf{h}(\mathbf{y}^{(j)})^T$.

III. Proposition 5 guarantees that for any j , $1 \leq j \leq \ell$, Bob can compute the error vector $\mathbf{e}^{(j)}$ and hence reconstruct $\mathbf{y}^{(j)} = \mathbf{x}^{(j)} + \mathbf{e}^{(j)}$ from his knowledge of $\mathbf{x}^{(j)}$. He can therefore open the mask $\mathbf{h}(\mathbf{y}^{(j)})^T$ and obtain the secret $s^{(j)}$.

The security of Protocol 6 can be proved by adapting the proof of Proposition 2 to the rank-metric case. Furthermore, Protocol 6 has polynomial cost in n, q and m for both computation and communication.

8 Concluding Remarks

We have presented a two-round PSMT protocol that has polynomial computational cost for both sender and receiver, and that achieves transmission rate linear in n and communication complexity in $O(n^2 \log n)$; we believe that our protocol is conceptually simpler compared to previous work and fully harnesses the properties of the pseudo-basis.

As proved in [20], the transfer rate is asymptotically optimal; furthermore, our protocol has a low multiplicative constant of 5. We moreover show that our vanilla protocol can be adapted to more general scenario, in the presence of a more powerful adversary.

It remains open whether the $O(n^2 \log n)$ communication complexity is optimal or not; the only known lower bound on this parameter is still $O(n)$, as the one for transfer rate [20]. It seems to us that a communication complexity lower than $O(n^2)$ is unlikely to be achievable, at least not without a completely different approach to the problem.

9 Acknowledgments

The authors would like to thank Serge Fehr, Ronald Cramer and Muriel Médard for their useful comments and suggestions.

References

- [1] S. Agarwal, R. Cramer, and R. de Haan. Asymptotically optimal two-round perfectly secure message transmission. In *Advances in Cryptology - CRYPTO 2006, 26th Annual International Cryptology Conference, Santa Barbara, California, USA, August 20-24, 2006, Proceedings*, pages 394–408, 2006.
- [2] R. Alshwede, Ning Cai, S.-Y.R. Li and R.W. Yeung. Network information flow. *IEEE Transactions on Information Theory*, volume 46, pages 1204–1216, Jul 2000.

- [3] Ph. Delsarte. Bilinear forms over a finite field, with applications to coding theory. *Journal of Combinatorial Theory, Series A*, volume 25, no. 3, pages 226–241, 1978.
- [4] D. Dolev, C. Dwork, O. Waarts, and M. Yung. Perfectly secure message transmission. *J. ACM*, 40(1):17–47, Jan. 1993.
- [5] È. M. Gabidulin. Theory of codes with maximum rank distance. *Problemy Peredachi Informatsii*, volume 21, pages 3–16, 1985.
- [6] J. Griggio. *Perfectly Secure Message Transmission Protocols with Low Communication Overhead and Their Generalization*. 2012. Master Thesis. <http://algant.eu/documents/theses/griggio.pdf>
- [7] S. Jaggi, M. Langberg, S. Katti, T. Ho, D. Katabi and M. Médard. Resilient Network Coding in the Presence of Byzantine Adversaries. *IEEE Transactions on Information Theory*, Volume 54, Issue 6, June 2008, pp. 2596 - 2603. Special Issue on Information-theoretic Security.
- [8] R. Koetter and M. Médard. An algebraic approach to network coding. *IEEE/ACM Transactions on Networking*, volume 11, no. 5, pages 782–795, Oct 2003.
- [9] K. Kurosawa and K. Suzuki. Truly efficient 2-round perfectly secure message transmission scheme. In *EUROCRYPT*, pages 324–340, 2008.
- [10] K. Kurosawa and K. Suzuki. Truly efficient 2 -round perfectly secure message transmission scheme. *Information Theory, IEEE Transactions on*, 55(11):5223–5232, Nov 2009.
- [11] S.-Y.R. Li, R.W. Yeung and Ning Cai. Linear network coding. *IEEE Transactions on Information Theory*, volume 49, pages 371–381, Feb 2003.
- [12] D. S. Lun, M. Médard, R. Koetter, M. Effros. On Coding for Reliable Communication over Packet Networks. *Physical Communication*, Volume 1, Issue 1, March 2008, pp. 3-2.
- [13] F. MacWilliams and N. Sloane. *The Theory of Error Correcting Codes*. North-Holland mathematical library. North-Holland Publishing Company, 1977.
- [14] J. L. Massey. Some applications of coding theory in cryptography. In *Codes and Ciphers: Cryptography and Coding IV*, pages 33–47, 1995.
- [15] H. M. Sayeed and H. Abu-Amara. Efficient perfectly secure message transmission in synchronous networks. *Information and Computation*, 126(1):53 – 61, 1996.
- [16] A. Shamir. How to Share a Secret. *Communications of the ACM*, 22(11):612-613, Nov 1979.
- [17] D. Silva and F. R. Kschischang. Security for wiretap networks via rank-metric codes. In *ISIT 2008, IEEE International Symposium on Information Theory*, pages 176–180, Jul 2008.
- [18] D. Silva and F. R. Kschischang. Universal Secure Network Coding via Rank-Metric Codes. *IEEE Transactions on Information Theory*, volume 57, pages 1124–1135, Feb 2011.
- [19] G. Spini and G. Zémor. Universally Secure Network Coding with Feedback. In *ISIT 2016, IEEE International Symposium on Information Theory*, Jul 2016.

- [20] K. Srinathan, A. Narayanan, and C. Rangan. Optimal perfectly secure message transmission. In M. Franklin, editor, *Advances in Cryptology – CRYPTO 2004*, volume 3152 of LNCS, pages 545–561. Springer Berlin Heidelberg, 2004.
- [21] K. Sundararajan, D. Shah, M. Médard, S. Jakubczak, M. Mitzenmacher and J. Barros. Network Coding Meets TCP: Theory and Implementation. *Proceedings of the IEEE*, invited paper, March 2011, pp. 490 – 512.